RIGHT-ANGLED ARTIN GROUPS WITH NON-PATH-CONNECTED BOUNDARY

WES CAMP

Abstract. We place conditions on the presentation graph Γ of a right-angled Artin group A_Γ that guarantee the standard CAT(0) cube complex on which A_Γ acts geometrically has non-path-connected boundary.

1. Introduction

In [7], Gromov showed that if G is a hyperbolic group acting geometrically on two metric spaces X and Y, then the boundaries of X and Y are homeomorphic. The same is not true for CAT(0) spaces; in [6] Croke and Kleiner demonstrate a group that acts geometrically on two CAT(0) spaces with non-homeomorphic boundaries, and it was later shown ([14]) that the same group has uncountably many distinct CAT(0) boundaries. The group is the right-angled Artin group whose presentation graph is the path on four vertices P_4, and so has presentation

$$\langle a, b, c, d \mid [a, b] = [b, c] = [c, d] = 1\rangle.$$

In [5], it is shown that the boundary of the standard CAT(0) cube complex on which this group acts is non-path-connected. The boundary of such a cube complex is connected if and only if the the presentation graph of the group is connected (and so the group is one-ended). In this paper, the method in [5] is generalized to a class of right-angled Artin groups whose presentation graphs admit a certain type of splitting. The main theorem here is as follows:

Theorem 1.1. *Let Γ be a connected graph. Suppose Γ contains an induced subgraph $(\{a, b, c, d\}, \{\{a, b\}, \{b, c\}, \{c, d\}\})$ (isomorphic to P_4), and there are subsets $B \subset lk(c)$ and $C \subset lk(b)$ with the following properties:*

1. *B separates c from a in Γ, with $d \notin B$;*
2. *C separates b from d in Γ, with $a \notin C$;*
3. *$B \cap C = \emptyset$.*

Then $\partial\mathcal{S}_\Gamma$ is not path connected.

Here, $\mathcal{S}_\Gamma$ is the standard CAT(0) cube complex on which the right-angled Artin group A_Γ with presentation graph Γ acts geometrically, and $lk(v)$ is the set of vertices of Γ sharing an edge with v. We in fact show a slightly stronger result, with the hypothesis $B \cap C = \emptyset$ replaced with the statement of Claim 3.7. The hypotheses here essentially require a copy of P_4 in Γ that is either not contained in a cycle, or has every cycle containing it separated by chords based at b and c. It is a known fact of graph theory that any graph that does not split as a join contains an induced subgraph isomorphic to P_4, and any graph Γ that splits as a non-trivial join has $\partial\mathcal{S}_\Gamma$ path connected, so the hypothesis that Γ contain a copy of P_4 is satisfied in any interesting case.

If a connected boundary of a CAT(0) space is locally connected, then it is a Peano space (a continuous image of [0,1]) and therefore path connected. The boundaries of some right-angled Coxeter groups are therefore known to be path connected ([11] and [4]), because they are locally connected. However, a consequence of a theorem in [10] is that for right-angled Artin groups, $\partial\mathcal{S}_\Gamma$ is locally connected iff Γ is a complete graph; i.e. $A_\Gamma \cong \mathbb{Z}^n$ and $\partial\mathcal{S}_\Gamma \cong S^{n-1}$. Thus no approach involving local connectivity works for right-angled Artin groups.

In [12], the construction of [6] is generalized to demonstrate a class of groups with non-unique boundary. These groups are of the form

$$G = (G_1 \times \mathbb{Z}^n) *_{\mathbb{Z}^n} (\mathbb{Z}^n \times \mathbb{Z}^m) *_{\mathbb{Z}^m} (\mathbb{Z}^m \times G_2),$$

where G_1 and G_2 are infinite CAT(0) groups. It is easily verified that if G_1 and G_2 are right-angled Artin groups, then G is a right-angled Artin group whose presentation graph satisfies the conditions of the main theorem of this paper; in fact, the method of this paper should work even if G_1 and G_2 are arbitrary infinite CAT(0) groups.

It seems this boundary path connectivity problem may be related to the question of when two right-angled Artin groups are quasi-isometric. In [1], Behrstock and Neumann show that all right-angled Artin groups whose presentation graphs are trees of diameter greater than 2 are quasi-isometric; in [2], Bestvina, Kleiner, and Sageev show that right-angled Artin groups with atomic presentation graphs (no valence 1 vertices, no separating vertex stars, and no cycles of length ≤ 4) have A_Γ quasi-isometric to $A_{\Gamma'}$ iff $\Gamma \cong \Gamma'$. The connection between these results and the result of this paper is that if Γ is a tree of diameter greater than 2, then Γ satisfies the hypotheses of the main theorem here, and therefore $\partial\mathcal{S}_\Gamma$ has non-path-connected boundary; if Γ is atomic, then Γ cannot satisfy the hypotheses of the main theorem here.

The author would like to thank Mike Mihalik for his guidance during the writing of this paper.

2. Preliminaries

Definition 2.1. Given a (undirected) graph Γ with vertex set $S = a_1, \ldots, a_n$, the corresponding **right-angled Artin group** A_Γ is the group with presentation

$$\langle a_1, \ldots, a_n \mid [a_i, a_j] \text{ if } i < j \text{ and } \{a_i, a_j\} \text{ is an edge of } \Gamma\rangle.$$

We call Γ the **presentation graph** for A_Γ.

Definition 2.2. If A_Γ is a right-angled Artin group with Cayley graph Λ_Γ, let $\overline{e} \in S$ be the label of the edge e of Λ_Γ. An **edge path** $\alpha \equiv (e_1, e_2, \ldots, e_n)$ in Λ_Γ is a map $\alpha : [0, n] \to \Lambda_\Gamma$ such that α maps $[i, i+1]$ isometrically to the edge e_i. For α an edge path in Λ_Γ, let $lett(\alpha) \equiv \{\overline{e}_1, \ldots, \overline{e}_n\}$, and let $\overline{\alpha} \equiv \overline{e}_1 \cdots \overline{e}_n$. If β is another geodesic with the same initial and terminal points as α, then call β a **rearrangement** of α.

Lemma 2.3. *If $w = g_1 \ldots g_k$ is a word in A_Γ (with each $g_i \in S^{\pm}$) that is not of minimal length, then two letters of $g_1 \ldots g_k$* ***delete****; that is, for some $i < j$, $g_i = g_j^{-1}$, the sets $\{g_i, g_j\}$ and $\{g_{i+1}, \ldots, g_{j-1}\}$ commute, and $w = g_1 \ldots g_{i-1} g_{i+1} \ldots g_{j-1} g_{j+1} \ldots g_k$.*

Proof. Let $w = h_1 \ldots h_m$ be a minimal length word representing w, and draw a van Kampen diagram D for the loop $g_1 \ldots g_k h_m^{-1} \ldots h_1^{-1}$. For each boundary edge

e_i corresponding to a g_i, trace a band across the diagram by picking the opposite edge of e_i in the relation square containing e_i, and continuing to pick opposite edges (without going backwards). Note that such a band cannot cross itself, and so this band must end on another boundary edge of D. Since $k > m$, there is some boundary edge e_i corresponding to some g_i that has its band B end on a boundary edge e_j corresponding to g_j, with $i < j$. Note this implies $g_i = g_j^{-1}$. Now, either all the bands corresponding to $g_{i+1}, \ldots, g_{j-1}$ cross B (implying each of $g_{i+1}, \ldots, g_{j-1}$ commutes with g_i and g_j), or some band corresponding to one of $g_{i+1}, \ldots, g_{j-1}$ ends on a boundary edge corresponding to another of $g_{i+1}, \ldots, g_{j-1}$. Picking an "innermost" such band and repeating the above argument gives the desired result. □

Remark 2.4. Note that the bands in the van Kampen diagram D share the same labels along their 'sides'. This means that deleting the band B from the diagram and matching up the separate parts of what remains (along paths with the same labels) gives a van Kampen diagram D' for the loop
$w = g_1 \ldots g_{i-1} g_{i+1} \ldots g_{j-1} g_{j+1} \ldots g_k h_m^{-1} \ldots h_1^{-1}$.

Remark 2.5. Given a non-geodesic edge path $(e_1, \ldots, e_k)$ in the Cayley graph Λ_Γ for A_Γ, we say edges e_i and e_j delete if their corresponding labels delete in the word $\overline{e_1} \ldots \overline{e_k}$.

Lemma 2.6. *Suppose A_Γ is a right-angled Artin group, and (α_1, α_2) and (β_1, β_2) are geodesics between the same two points in in the Cayley graph Λ_Γ for A_Γ. There exist geodesics $(\gamma_1, \tau_1), (\gamma_1, \delta_1), (\delta_2, \gamma_2)$, and (τ_2, γ_2) with the same end points as $\alpha_1, \beta_1, \alpha_2, \beta_2$ respectively, such that:*

1. *τ_1 and τ_2 have the same labels,*
2. *δ_1 and δ_2 have the same labels, and*
3. *$lett(\tau_1)$ and $lett(\delta_1)$ are disjoint and commute.*

Furthermore, the paths (τ_1^{-1}, δ_1) and (δ_2, τ_2^{-1}) are geodesic.

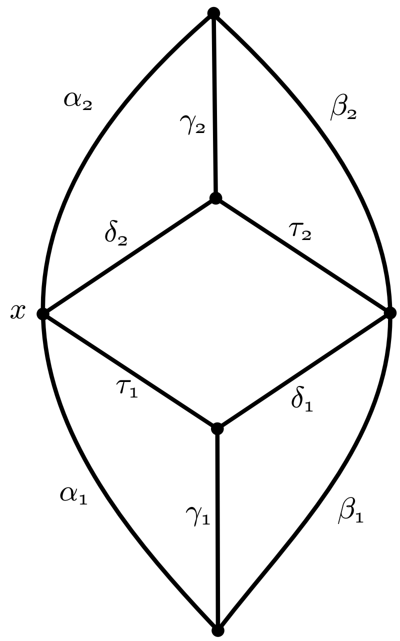

FIGURE 1

Proof. Let D be a van Kampen diagram for the loop $(\alpha_1, \alpha_2, \beta_2^{-1}, \beta_1^{-1})$, and let $\alpha_1 = (a_1, \ldots, a_k)$, $\beta_1 = (b_1, \ldots, b_m)$. Let $a_{i_1}, \ldots, a_{i_j}$ be (in order) the edges of α_1 whose bands in D end on β_1. Note that by Lemma 2.3, β_1 can be rearranged to begin with an edge labeled $\overline{a_{i_1}}$, since a_{i_1} and b_{ℓ_1} delete in (α_1^{-1}, β_1) for some ℓ_1 and all the bands based at $b_1, \ldots, b_\ell, a_1, \ldots, a_{i_1-1}$ cross the band based at a_{i_1} and ending at b_{ℓ_1}. Similarly, β_1 can be rearranged to begin with an edge labeled $\overline{a_{i_1}}$ followed by an edge labeled $\overline{a_{i_2}}$, and continuing in this manner, we obtain a rearrangement of β_1 that begins with $\gamma_1 = (\overline{a_{i_1}}, \ldots, \overline{a_{i_j}})$, and we let δ_1 be the remainder of this rearrangment. This argument also implies α_1 can be rearranged to begin with γ_1, and we let τ_1 be the remainder of this rearrangement. Note that if e is an edge of τ_1, no edge of δ_1 is labeled $\overline{e}$ or $\overline{e}^{-1}$, since bands with those labels must have crossed in D. We obtain γ_2, τ_2 and δ_2 in the analogous way from α_2 and β_2, and note that in a van Kampen diagram B' for $(\tau_1, \delta_2, \tau_2^{-1}, \delta_1^{-1})$, no band based on τ_1 can end on δ_2, since (τ_1, δ_2) is geodesic, and no band based on τ_1 ends on δ_1, since τ_1 and δ_1 share no labels or inverse labels. Therefore all bands on τ_1 end on τ_2, so τ_1 and τ_2 have the same labels, as do δ_1 and δ_2. $\square$

Definition 2.7. Under the hypotheses of the previous lemma, we call τ_1 the **down edge path** at x, and we call δ_2 the **up edge path** at x. If α_1 and β_1 have the same length, we call the above figure the **diamond** at x for (α_1, α_2) and (β_1, β_2).

Definition 2.8. P_4 is the (undirected) graph on four vertices a, b, c, d, with edge set $\{\{a, b\}, \{b, c\}, \{c, d\}\}$.

Definition 2.9. The **union** of two graphs (V_1, E_1) and (V_2, E_2) is the graph $(V_1 \cup V_2, E_1 \cup E_2)$.

Definition 2.10. The **join** of two graphs (V_1, E_1) and (V_2, E_2) is the graph $(V_1 \cup V_2, E_1 \cup E_2 \cup (V_1 \times V_2))$.

Definition 2.11. A graph is **decomposable** if it can be expressed as joins and unions of isolated vertices.

The following is Theorem 9.2 in [9].

Theorem 2.12. *A finite graph G is decomposable iff it does not contain P_4 as an induced subgraph.*

In particular, if a connected graph G does not contain P_4 as an induced subgraph, then it must split as the join $G_1 \vee G_2$, for some subgraphs G_1, G_2 of G.

Definition 2.13. For a graph Γ and a vertex a of Γ, $lk(a) = \{b \in \Gamma \mid \{a, b\}$ is an edge of $\Gamma\}$.

Let Λ_Γ be the Cayley graph for the group A_Γ.

Definition 2.14. The **standard complex** $\mathcal{S}_\Gamma$ for the group A_Γ is the CAT(0) cube complex whose one-skeleton is Λ_Γ, with each cube given the geometry of $[0,1]^n$ for the appropriate n.

For more on cube complexes and the definitions below, see [13].

Definition 2.15. A **midcube** in a cube complex C is the codimension 1 subspace of an n-cube $[0,1]^n$ obtained by restricting exactly one coordinate to $\frac{1}{2}$. A **hyperplane** is a connected nonempty subspace of C whose intersection with each cube is either empty or consists of one of its midcubes.

Lemma 2.16. *If D is a hyperplane of the cube complex C, then $C - D$ has exactly two components.*

Given a graph Γ, a vertex v of Γ, and the corresponding standard complex $\mathcal{S}_\Gamma$, note that if a hyperplane of $\mathcal{S}_\Gamma$ intersects an edge of $\mathcal{S}_\Gamma$ with label v, then every edge intersected by this hyperplane is also labeled v. Thus we can refer to hyperplanes in $\mathcal{S}_\Gamma$ as v-hyperplanes, for v a vertex of Γ. If x is a vertex of $\mathcal{S}_\Gamma$, then xv and x are separated by a v-hyperplane D. Let $x\mathcal{S}_{lk(v)}$ denote the cube complex generated by the coset $x\langle lk(v)\rangle$; then D and $x\mathcal{S}_{lk(v)}$ are isometric and parallel, of distance $\frac{1}{2}$ apart.

Definition 2.17. A metric space (X, d) is **proper** if each closed ball is compact.

Definition 2.18. Let (X, d) be a proper CAT(0) space. Two geodesic rays $c, c' : [0, \infty) \to X$ are called **asymptotic** if for some constant K, $d(c(t), c'(t)) \leq K$ for all $t \in [0, \infty)$. Clearly this is an equivalence relation on all geodesic rays in X. We define the **boundary** of X (denoted ∂X) to be the set of equivalence classes of geodesic rays in X. We denote the union $X \cup \partial X$ by $\overline{X}$.

The next proposition guarantees that the topology we wish to put on the boundary is independent of our choice of basepoint in X.

Proposition 2.19. *Let (X, d) be a proper CAT(0) space, and let $c : [0, \infty) \to X$ be a geodesic ray. For a given point $x \in X$, there is a unique geodesic ray based at x which is asymptotic to c.*

For a proof of this (and more details on what follows), see [3].

We wish to define a topology on $\overline{X}$ that induces the metric topology on X. Given a point in ∂X, we define a neighborhood basis for the point as follows:
Pick a basepoint $x_0 \in X$. Let c be a geodesic ray starting at x_0, and let $\epsilon > 0$, $r > 0$. Let $S(x_0, r)$ denote the sphere of radius r centered at x_0, let $B(x_0, r)$ denote

the open ball of radius r centered at x_0 and let $p_r : X - B(x_0, r) \to S(x_0, r)$ denote the projection to $S(x_0, r)$. Define

$$U(c, r, \epsilon) = \{x \in \overline{X} : d(x, x_0) > r, d(p_r(x), c(r)) < \epsilon\}.$$

This consists of all points in $\overline{X}$ whose projection to $S(x_0, r)$ is within ϵ of the point of the sphere through which c passes. These sets together with the metric balls in X form a basis for the **cone topology**. The set ∂X with this topology is sometimes called the **visual boundary**. In this article, we will call it the boundary of X.

Proposition 2.20. *If X and Y are proper CAT(0) spaces, then $\partial(X \times Y) \cong \partial X * \partial Y$, where $*$ denotes the spherical join.*

If the graph Γ splits as a non-trivial join $\Gamma_1 \vee \Gamma_2$, then the group A_Γ splits as the direct product $A_{\Gamma_1} \times A_{\Gamma_2}$, and so we have $\mathcal{S}_\Gamma \cong \mathcal{S}_{\Gamma_1} \times \mathcal{S}_{\Gamma_2}$. The previous proposition then gives that $\partial\mathcal{S}_\Gamma \cong \partial\mathcal{S}_{\Gamma_1} * \partial\mathcal{S}_{\Gamma_2}$. Any non-trivial spherical join is path connected, and so $\partial\mathcal{S}_\Gamma$ is path connected.

Lemma 2.21. *There is a bound $\delta > 0$ such that if α is a CAT(0) geodesic path in $\mathcal{S}_\Gamma$, then there is a Cayley graph geodesic path β in Λ_Γ (contained naturally in $\mathcal{S}_\Gamma$) such that each vertex of β is within distance δ of α, and each point of α is within δ of a vertex of β.*

A proof of this can be found in Section 3 of [8].

3. Result

The goal of this section is to prove the following theorem:

Theorem 3.1. *Let Γ be a connected graph. Suppose Γ contains an induced subgraph $(\{a, b, c, d\}, \{\{a, b\}, \{b, c\}, \{c, d\}\})$ (isomorphic to P_4), and there are subsets $B \subset lk(c)$ and $C \subset lk(b)$ with the following properties:*

1. *B separates c from a in Γ, with $d \notin B$;*
2. *C separates b from d in Γ, with $a \notin C$;*
3. *$B \cap C = \emptyset$.*

Then $\partial\mathcal{S}_\Gamma$ is not path connected.

In fact, we prove a stronger result, with the hypothesis $B \cap C = \emptyset$ replaced by the statement of Claim 3.7. For the remainder of this section, suppose $a, b, c, d \in \Gamma$, $B \subset lk(c)$, and $C \subset lk(b)$ are as in Theorem 3.1. Note that $b \in B$, $c \in C$. We wish to consider the following rays in Λ_Γ (equivalently in $\mathcal{S}_\Gamma$), based at the identity vertex $*$:

$$r = cdab(cb)^2cdab(cb)^6 \cdots = \prod_{i=1}^{\infty}(cb)^{k_i}cdab$$

and

$$s = dbcb^2adbc(b^2c)^2b^2adbc(b^2c)^6b^2a \cdots = \prod_{i=1}^{\infty} dbc(b^2c)^{k_i}b^2a$$

where the k_i are defined recursively with $k_0 = -1$, $k_{i+1} = 2k_i + 2$.

Define the following vertices of r, for $n \geq 0$:

$$v_n = \left(\prod_{i=1}^{n}(cb)^{k_i}cdab\right)(cb)^{k_{n+1}}cd$$

$$v'_n = v_n a$$

Define the following vertices of s, for $n \geq 0$:

$$w_n = \left(\prod_{i=1}^{n} dbc(b^2c)^{k_i}b^2a\right)$$

$$w'_n = w_n d$$

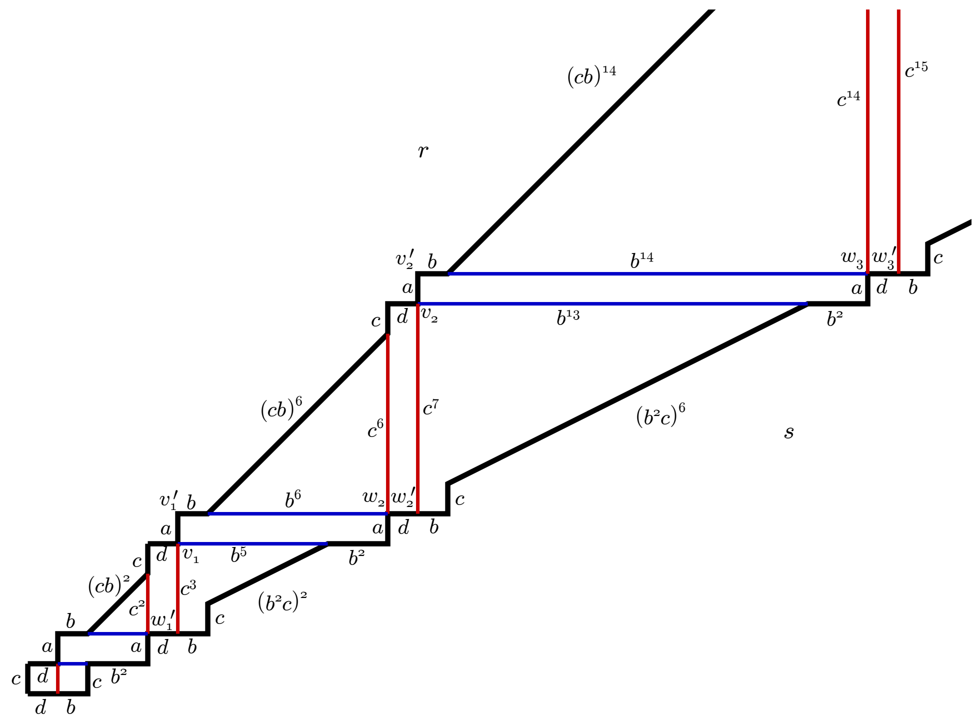

FIGURE 2

We have $v_0 = cd$, $v'_0 = cda$, $v_1 = cdab(cb)^2cd$, $w_0 = *$, $w'_0 = d$, $w_1 = dbcb^2a$. It will be helpful to refer to Figure 2 for many of the claims that follow.

The following is proved in [5].

Claim 3.2. *For $n \geq 0$, $v_n = w'_n c^{k_{n+1}+1}$ and $v'_n b^{k_{n+2}+1} = w_{n+1}$.*

Since $b \in B$ and $c \in C$, we then have $v_n\langle C\rangle = w'_n\langle C\rangle$ and $w_n\langle B\rangle = v'_{n-1}\langle B\rangle$.

If Q_c denotes the component of c in $\Gamma - B$, and Q_b denotes the component of b in $\Gamma - C$, then A_Γ can be represented as $\langle Q_c \cup B\rangle *_B \langle\Gamma - Q_c\rangle$ or $\langle Q_b \cup C\rangle *_C \langle\Gamma - Q_b\rangle$, and so at each vertex x of Λ_Γ, the cosets $x\langle B\rangle$ and $x\langle C\rangle$ separate Λ_Γ. Therefore, if $x\mathcal{S}_B$ and $x\mathcal{S}_C$ denote the cube complexes generated by $\langle B\rangle$ and $\langle C\rangle$ respectively at a vertex x of $\mathcal{S}_\Gamma$, then $x\mathcal{S}_B$ and $x\mathcal{S}_C$ separate $\mathcal{S}_\Gamma$. Note that $\mathcal{S}_\Gamma - x\mathcal{S}_B$ has at least

two components: one containing xc^{-1}, and one containing xa. Similarly, $\mathcal{S}_\Gamma - x\mathcal{S}_C$ has at least two components: one containing xb^{-1}, and one containing xd.

For each i, define the following components of $\mathcal{S}_\Gamma$:

(1) V_i^+ is the component of $\mathcal{S}_\Gamma - v_i\mathcal{S}_B$ containing $v_i a$;
(2) V_i^- is the component of $\mathcal{S}_\Gamma - v_i\mathcal{S}_B$ containing $v_i c^{-1}$;
(3) W_i^+ is the component of $\mathcal{S}_\Gamma - w_i\mathcal{S}_C$ containing $w_i d$;
(4) W_i^- is the component of $\mathcal{S}_\Gamma - w_i\mathcal{S}_C$ containing $w_i b^{-1}$.

Note V_i^+ contains the vertices of r after v_i, and W_i^+ contains the vertices of s after w_i. For each $V_i^\pm$, (respectively $W_i^\pm$), let $\overline{V_i^\pm}$ denote the closure of $V_i^\pm$ in $\mathcal{S}_\Gamma$, so $\overline{V_i^\pm} = V_i^\pm \cup v_i\mathcal{S}_B$ ($\overline{W_i^\pm} = W_i^\pm \cup w_i\mathcal{S}_C$). For a subset S of $\mathcal{S}_\Gamma$, let $L(S)$ denote the limit set of S in $\partial\mathcal{S}_\Gamma$.

Claim 3.3. (1) *The sets $\overline{V_i^\pm}$, $\overline{W_i^\pm}$ are convex.*
(2) *$L(\overline{V_i^+}) \cap L(\overline{V_i^-}) = L(v_i\mathcal{S}_B)$ and $L(\overline{W_i^+}) \cap L(\overline{W_i^-}) = L(w_i\mathcal{S}_C)$.*
(3) *The set $L(v_i\mathcal{S}_B)$ (respectively $L(w_i\mathcal{S}_C)$) separates $L(\overline{V_i^+})$ and $L(\overline{V_i^-})$ (respectively $L(\overline{W_i^+})$ and $L(\overline{W_i^-})$) in ∂X.*

Proof. For (1), the only way out of the set $\overline{V_i^+}$ is through the convex subcomplex $v_i\mathcal{S}_B$.

For (2), if q is a ray in $L(\overline{V_i^+}) \cap L(\overline{V_i^-})$, then there are geodesic rays $q_1 \in \overline{V_i^+}$, $q_2 \in \overline{V_i^-}$ that are a bounded distance from q, and therefore from one another. Thus both q_1 and q_2 remain a bounded distance from $v_i\mathcal{S}_B$, as required.

For (3), suppose $\alpha : [0,1] \to \partial\mathcal{S}_\Gamma$ is a path connecting $x \in L(V_i^+)$ and $y \in L(V_i^-)$. Choose $w \in v_i\mathcal{S}_B$, and for each $t \in [0,1]$, let $\beta_t : [0,\infty) \to \mathcal{S}_\Gamma$ be the geodesic ray from w to $\alpha(t) \in \partial\mathcal{S}_\Gamma$. This gives a continuous map $H : [0,1]\times[0,\infty) \to \mathcal{S}_\Gamma$ where $H(t,s) = \beta_t(s)$. Note $H(0,s) \subset V_i^+$, $H(1,s) \subset V_i^-$. For each $n \geq 0$, let z_n be a point of $H([0,1]\times\{n\})$ in $v_i\mathcal{S}_B$; then $L(\cup_{n=1}^\infty\{z_n\}) \subset Im(\alpha) \cap L(v_i\mathcal{S}_B)$ as required. □

In [5], it is shown that r and s track distinct CAT(0) geodesics in $\mathcal{S}_\Gamma$, so $L(r)$ and $L(s)$ are distinct one-element sets.

Claim 3.4. *For $n \geq 1$, the sets $L(w_{2n-1}\mathcal{S}_C)$ and $L(r)$ are separated in $\partial\mathcal{S}_\Gamma$ by $L(v_{2n+1}\mathcal{S}_B)$.*

Proof. First note that $L(r) \in L(V_i^+)$ for each $i \geq 1$. Let D_{2n} be the d-hyperplane that separates w_{2n} from w'_{2n} (and also separates v_{2n} from the previous vertex of r), and let A_{2n} be the a-hyperplane that separates v_{2n} from v'_{2n} (and also separates w_{2n+1} from the previous vertex of s). Note that $w_{2n-1}\mathcal{S}_C$ is contained in the same component of $\mathcal{S}_\Gamma - D_{2n}$ as $*$ since $d \notin C$ and therefore no path in $\langle C\rangle$ based at w_{2n-1} crosses D_{2n}. Also note $A_{2n} \subset V_{2n+1}^-$. Since D_{2n} and A_{2n} cannot cross (since d does not commute with a), and D_{2n} is not in the same component as $v_{2n+1}\mathcal{S}_B$ in $\mathcal{S}_\Gamma - A_{2n}$, we have that $w_{2n-1}\mathcal{S}_C \subset V_{2n+1}^-$. The previous claim gives the result. □

Claim 3.5. *For $n \geq 1$, the sets $L(v_{2n-1}\mathcal{S}_B)$ and $L(r)$ are separated in $\partial\mathcal{S}_\Gamma$ by $L(w_{2n+1}\mathcal{S}_C)$.*

Proof. The proof is analagous to the proof of the previous claim, replacing the hyperplanes D_{2n} and A_{2n} with the hyperplanes A_{2n-1} and D_{2n} respectively. □

Remark 3.6. The previous two claims imply that if there is a path in $\partial\mathcal{S}_\Gamma$ between a point of $L(w_1\mathcal{S}_C)$ and $L(r)$, the path must pass through (in order) $L(v_3\mathcal{S}_B)$, $L(w_5\mathcal{S}_C)$, $L(v_7\mathcal{S}_B)$, $L(w_9\mathcal{S}_C)$, and so on.

We will now show that the sets $L(v_i\mathcal{S}_B)$ (resp. $L(w_i\mathcal{S}_C)$) are eventually 'close' to $L(s)$ (resp. $L(r)$), implying the path described in Remark 3.6 cannot exist.

Claim 3.7. *$C \cap lk(a) \cap lk(d) = C \cap lk(a) \cap lk(c) = \emptyset$, and $B \cap lk(a) \cap lk(d) = B \cap lk(d) \cap lk(b) = \emptyset$.*

Proof. If $e \in C \cap lk(a) \cap lk(d)$, then (a, e, d, c) is a path from a to c in Γ. Since B separates a from c and $d \notin B$, we must have $e \in B$, but $B \cap C = \emptyset$. Similarly, if $e \in C \cap lk(a) \cap lk(c)$, then (a, e, c) is a path from a to c in Γ, and so $e \in B$, contradiction. The remaining statements are proved identically. □

For $i \geq 1$, let r_i (respectively s_i) be the segment of r (respectively s) between $*$ and v_i' (respectively $*$ and w_i'). Let β_i be a Cayley graph geodesic ray based at w_i' with labels in B, and let γ_i be a Cayley graph geodesic ray based at v_i' with labels in C.

Claim 3.8. *Any Cayley graph geodesic from $*$ to a point of γ_i must pass within 4 units of v_i'. Any Cayley graph geodesic from $*$ to a point of β_i must pass within 4 units of w_i'.*

Proof. First observe that if (r_i, γ_i) is not Λ_Γ-geodesic, then an edge of γ_i must delete with an edge of r_i. Since $a, b, d \notin C$, the labels of these deleting edges must be c and c^{-1}. However, the labels of these edges must also be in $lk(a) \cap lk(d)$, by Lemma 2.3 (see Figure 2). Therefore (r_i, γ_i) is a Cayley geodesic.

Now, suppose there is a Λ_Γ-geodesic ρ between $*$ and a point of γ_i with $d(\rho, v_i') > 4$. Let α denote the segment of (r_i, γ_i) between $*$ and the endpoint of ρ. Consider a diamond based at v_i' for ρ and α as in Lemma 2.6. Let τ and δ be the down edge path and up edge path respectively at v_i', and note τ and δ have length at least 3. Every Λ_Γ-geodesic from $*$ to v_i' must end with an edge labeled a, so every label of δ is in $lk(a)$. If an edge of τ has label d, then every label of δ is in $C \cap lk(a) \cap lk(d)$, but this set is empty by Claim 3.7. By Lemma 2.3 every other edge of τ has its label in $lk(d) \cap \{a, b, c, d\}$, so the remaining edges of τ must be labeled c, but $C \cap lk(a) \cap lk(c)$ is also empty. Thus $d(\rho, v_i') \leq 4$. The proof of the second statement is identical. □

Claim 3.9. *$\partial\mathcal{S}_\Gamma$ is not path connected.*

Proof. Observe that since $v_{n-1}' b^{k_{n+1}+1} = w_n$ by Claim 3.2 and $C \subset lk(b)$, any ray α based at w_n with labels in C stays a bounded distance from the ray based at v_{n-1}' with the same labels. Combining Claim 3.8 and Lemma 2.21, we have that a CAT(0) geodesic from $*$ to a point of $L(\alpha)$ must pass within $\delta + 4$ of v_{n-1}', where δ is the tracking constant given by Lemma 2.21. We therefore have that any sequence of points $\{p_i\}_{i=1}^\infty$ with each $p_i \in L(w_i\mathcal{S}_C) \subset \partial\mathcal{S}_\Gamma$ must converge to $L(r) \in \partial\mathcal{S}_\Gamma$. Similarly, any sequence of points $\{q_i\}_{i=1}^\infty$ with each $q_i \in L(v_i\mathcal{S}_B) \subset \partial\mathcal{S}_\Gamma$ must converge to $L(s) \in \partial\mathcal{S}_\Gamma$. Therefore, by Remark 3.6, given any ϵ, any path from a point of $L(w_1\mathcal{S}_C)$ to $L(r)$ eventually bounces back and forth infinitely between the ϵ-neighborhood of $L(s)$ and the ϵ-neighborhood of $L(r)$, which is impossible; therefore, no such path exists. □

References

[1] J. Behrstock and W. Neumann. *Quasi-isometric classification of graph manifold groups.* Duke Math. J. 141 (2008), 217-240.

[2] M. Bestvina, B. Kleiner, and M. Sageev. *The asymptotic geometry of right-angled Artin groups, I.* Geom. Topol. 12 (2008), 1653-1699.

[3] M.R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature.* Grundlehren Math. Wissensch. 319 (Springer-Verlag, 1999).

[4] W. Camp and M. Mihalik. *A classification of right-angled Coxeter groups with no 3-flats and locally connected boundary.* Submitted for publication. arXiv:1206.5234 [math.GR], 2012.

[5] G. Conner, M. Mihalik and S. Tschantz. *Homotopy of Ends and Boundaries of CAT(0) Groups.* Geom. Dedicata 120 (1) (2006), 1-17.

[6] C. Croke and B. Kleiner. *Spaces with non-positive curvature and their ideal boundaries.* Topology 39 (2000), 549-556.

[7] M. Gromov. *Hyperbolic groups.* Essays in Group Theory, S. Gersten ed., MSRI Publications 8 (1987), 75-265.

[8] T. Hsu and D. Wise. *Separating quasiconvex subgroups of right-angled Artin groups.* Math. Z. 240 (2002), 521-548.

[9] R. Meris. *Graph Theory.* Wiley, New York, 2001.

[10] M. Mihalik and K. Ruane. *CAT(0) groups with non-locally connected boundary.* J. London Math. Soc. (2) 60 (1999), 757-770.

[11] M. Mihalik, K. Ruane, and S. Tschantz. *Local connectivity of right-angled Coxeter group boundaries.* J. Group Theory 10 (2007), 531-560.

[12] C. Mooney. *Generalizing the Croke-Kleiner construction.* Topology Appl. 157 (7) (2010), 1168-1181.

[13] Michah Sageev. *Ends of group pairs and non-positively curved cube complexes.* Proc. London Math. Soc. (3) 71 (1995), 585-617.

[14] J. Wilson. *A CAT(0) group with uncountably many distinct boundaries.* J. Group Theory 8 (2005), 229-238.

Department of Mathematics, Vanderbilt University, Nashville, TN 37240
E-mail address: w.camp@vanderbilt.edu